

George Palikaras
Director | Co-founder
BlueLedger AI Inc.
george@blueledger.ai

1 902 222 4423

April 17, 2026

Ms. Vanessa A. Countryman
Secretary
U.S. Securities and Exchange Commission
100 F Street, N.E.
Washington, D.C. 20549-1090

Via: rule-comments@sec.gov and the Commission's public comment portal

Re: Public Comment on the Consolidated Audit Trail Data Storage Amendment — File No. 4-698; Release Nos. 34-104504 (Dec. 23, 2025), 34-105003, and 34-105107

Dear Ms. Countryman:

I respectfully submit this comment in response to the Commission's action on the 2025 Cost Savings Amendment and the related Data Storage Amendment to the National Market System Plan Governing the Consolidated Audit Trail (the "CAT NMS Plan"). The submission is offered in the spirit of the Commission's statutory mandate to protect investors, maintain fair and orderly markets, and facilitate capital formation — and in support of the Commission's stated interest in reducing CAT operating costs without compromising its regulatory utility.

The core finding of this submission is that the cost savings associated with shortening CAT data retention from five years to three years (itself a step-down from the original six-year online-retention requirement of the CAT NMS Plan, reduced to five years by the 2025 Cost Savings Exemptive Order) are modest in absolute terms and modestly lower still when measured against the expanded cloud-archival options available in 2026. The corresponding loss of regulatory capability, by contrast, is structural. Complex cross-market manipulation cases — from the 2010 Flash Crash prosecution of Navinder Sarao to the 2020 JPMorgan spoofing settlement to the 2024 Archegos convictions — routinely mature on enforcement timelines of five to ten years. A three-year CAT, in practical effect, renders the most consequential investigations dependent on fragmented broker-dealer records precisely when the Commission needs consolidated order-flow reconstruction most.

This comment demonstrates that a seven-year retention horizon is not only consistent with recordkeeping obligations under Rule 17a-4, Sarbanes-Oxley, the Commodity Exchange Act, the Bank Secrecy Act, and

MiFID II — but can be implemented at lower total cost than the approved three-year architecture through a properly designed tiered cloud storage policy (S3 Standard → Intelligent-Tiering → Glacier Instant Retrieval → Glacier Deep Archive) combined with a federal enterprise discount agreement. The submission concludes with proposed redline language for Section 6.5(b)(i) of the CAT NMS Plan and the related data-availability provisions of Appendix D.

I appreciate the Commission's consideration of this comment and stand ready to provide supporting data, engineering documentation, or additional analysis to the Division of Trading and Markets, the Division of Economic and Risk Analysis, and the CAT NMS Plan Operating Committee.

Respectfully submitted,



George Palikaras

Director | Co-founder

BlueLedger AI

Enclosures:

1. Main Submission;
2. Exhibits A–E (Comparative Retention Matrix, Enforcement Timeline, AWS Cost-Tier Analysis, Three-Scenario Cost Model, Redline of Section 6.5(b)(i) and Appendix D § 6.4).

PUBLIC COMMENT SUBMISSION

Consolidated Audit Trail Data Retention and Cost Optimization

File No. 4-698

Release Nos. 34-104504, 34-105003, 34-105107

Executive Summary

The 2025 Cost Savings Amendment, as modified and approved by the Commission on March 27, 2026 in Release No. 34-105107, reduces Consolidated Audit Trail ("CAT") data retention from five years to three years — the culmination of a reduction series that began with the CAT NMS Plan's original six-year online-retention requirement and proceeded through the September 2025 Cost Savings Exemptive Order's interim reduction to five years. The Data Storage Amendment is projected in Release 34-104504 to save approximately \$23.5 to \$32 million annually in cloud-hosting costs, of which approximately \$6.5 to \$9 million represents the marginal increase attributable to the Amendment's extension of the September 2025 Exemptive Order (reducing retention from five to three years, OMM Quotes from one year to six months, and Options SIP Data from five years to six months). The broader seven-component amendment is projected by the Commission to reduce total annual CAT operating costs by approximately \$50 to \$70 million against a 2025 operating budget of approximately \$228 million (as revised in May 2025 from an initial \$249 million). The retention-specific portion of these savings is the focus of this comment.

The central conclusion of this submission is that the Commission can deliver greater cost savings than the Data Storage Amendment while simultaneously extending retention to seven years, by substituting a tiered cloud-archival architecture (Glacier Deep Archive with Object Lock) for the current approach and renegotiating the federal enterprise discount agreement with Amazon Web Services. Retaining CAT data for seven years is cheaper, more defensible, and better aligned with the Commission's statutory remedies than deleting it after three.

The three strongest factual anchors

1. Enforcement timelines structurally exceed three years. Every major recent market-manipulation case cited in this comment — Sarao (5 years to charges, 10 to sentencing), JPMorgan spoofing (up to 12 years from earliest conduct to settlement), 1MDB (7–8 years to SEC resolution), Archegos (3.5 years to sentencing), SAC/Martoma (6 years to conviction) — required consolidated order-flow or related-market data well beyond a three-year horizon.
2. The Commission's own remedies presuppose longer data. The National Defense Authorization Act of 2021 (NDAA § 6501) codified a 10-year statute of limitations for scienter-based disgorgement under Exchange Act § 10(b), Securities Act § 17(a)(1), and Investment Advisers Act § 206(1). A three-year CAT retention window cannot support a 10-year disgorgement remedy. The misalignment is a policy inconsistency, not merely a cost question.

3. The cost argument is inverted by 2026 cloud economics. The ratio of S3 Standard (\$0.023/GB-month) to S3 Glacier Deep Archive (\$0.00099/GB-month) is approximately 23×. Storing aged CAT data in Deep Archive with S3 Object Lock costs on the order of \$11,880 per petabyte per year at list, and materially less with a federal enterprise discount agreement. Extending retention to seven years under a tiered architecture is achievable at a lower annual cost than the approved three-year deletion amendment.

The core ask

- (1) Extend core CAT order and execution retention to seven years. Amend Section 6.5(b)(i) of the CAT NMS Plan to replace the "three (3) years" retention horizon with "seven (7) years" for core order lifecycle and execution records, while preserving shorter horizons for specified categories of raw personally identifiable information.
- (2) Adopt a tiered cloud-archival storage model. Replace the current storage approach with an explicit tier lifecycle: S3 Standard (days 0–90), S3 Intelligent-Tiering (days 91–365), S3 Glacier Instant Retrieval (years 1–3), and S3 Glacier Deep Archive with Object Lock (years 3–7), combined with columnar re-encoding (Parquet + Zstandard).
- (3) Direct CAT LLC to renegotiate a federal enterprise discount agreement with AWS and solicit competitive bids from Microsoft Azure Government and Google Cloud Assured Workloads, with zero-egress provisions for regulatory pulls by the SEC and SROs.
- (4) Codify explicit redline language in Section 6.5(b)(i) of the CAT NMS Plan and the related provisions of Appendix D, as set forth in Section IX below, to reduce the Commission's drafting burden and align the CAT NMS Plan with Rule 17a-4, Sarbanes-Oxley § 802, CFTC Rule 1.31, the Bank Secrecy Act, and MiFID II.

I. Introduction and Standing

The Consolidated Audit Trail is the most significant market-surveillance infrastructure ever built in the United States. Adopted under Rule 613 of Regulation NMS following the 2010 Flash Crash, it consolidates order lifecycle data — every order, cancellation, modification, and execution — across U.S. equities and listed options markets into a single queryable repository. The Commission itself has described CAT data as foundational to modern cross-market surveillance; the Division of Enforcement publicly credited CAT with uncovering the Nuveen/Billimek front-running scheme in late 2022, a matter that produced approximately \$47 million in illegal profits and that several commentators concluded would have been difficult to detect absent consolidated cross-market visibility.

On December 18, 2025, CAT LLC filed a joint-industry proposed amendment to the CAT NMS Plan intended to reduce operating costs (the "2025 Cost Savings Amendment"). On December 23, 2025, the Commission published the filing as Release No. 34-104504. The broader amendment package was subsequently approved with modifications on March 27, 2026 in Release No. 34-105107, and the data retention component — the "Data Storage Amendment" — now permits deletion of CAT data older than three years (down from five, following the September 2025 Exemptive Order), OMM Quotes older than six months (down from one year), and Options SIP Data older than six months (down from five years). The Commission has estimated that these storage-specific changes in the Data Storage Amendment deliver approximately \$23.5 to \$32 million in total annual cloud-hosting savings, of which approximately \$6.5 to \$9 million represents the incremental savings from extending the September 2025 Exemptive Order.

This comment does not oppose the objective of cost reduction. It argues, to the contrary, that cost reduction and retention expansion are complementary objectives under the cloud-storage technology available in 2026. The three-year horizon is neither statutorily required nor technologically compelled. It is a choice — and it is the wrong choice for a regulatory asset whose value is measured in the detection of multi-year manipulation schemes that mature long after a three-year tape has been purged.

The submission is organized as follows. Section II sets out the empirical enforcement-timeline reality that makes a three-year horizon structurally insufficient. Section III benchmarks the CAT against international and adjacent-regime recordkeeping regimes. Section IV demonstrates alignment with existing U.S. statutory retention obligations. Section V addresses the cost argument directly, presenting a three-scenario model of CAT storage economics under current AWS pricing. Section VI outlines the proposed technical architecture. Section VII presents a case study illustrating the harm to post-event investigations under truncated retention. Section VIII discusses the regulatory-arbitrage risk created by a shortened window. Section IX offers specific redline amendments to Appendix D of the CAT NMS Plan. Section X concludes.

II. The Enforcement-Timeline Reality: Why Three Years Is Structurally Insufficient

The Commission's most consequential market-manipulation and fraud matters are not resolved within three years of the underlying conduct. They are resolved on timelines of five to twelve years, and the Commission's statutory remedies — particularly the 10-year disgorgement period codified in Section 6501 of the National Defense Authorization Act for Fiscal Year 2021 — explicitly contemplate lookback periods well beyond any three-year data horizon. The misalignment between the Commission's retention policy and its remedial authorities is the single most important structural point in this comment.

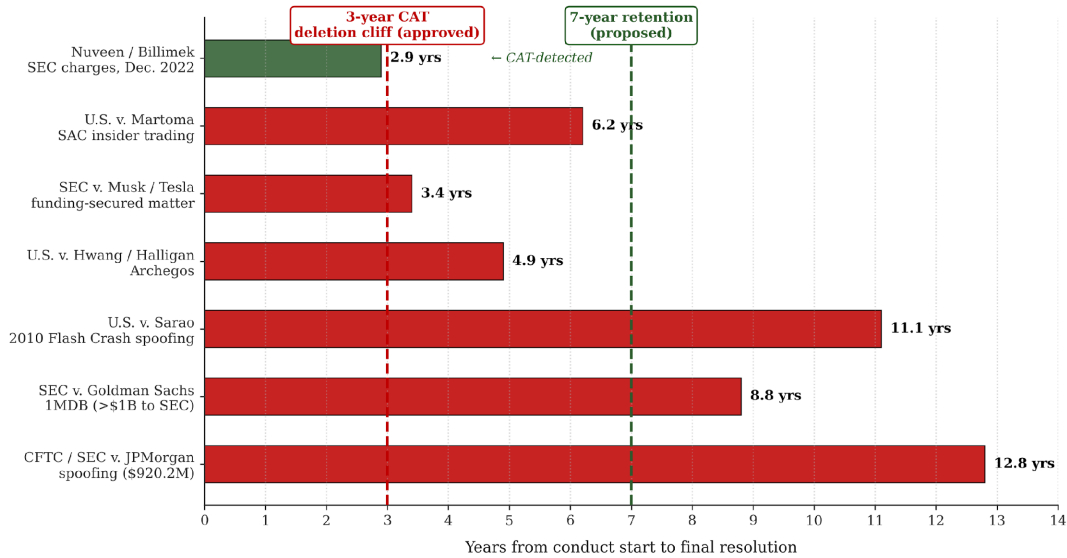
A. A representative enforcement-timeline record

The following matters, all publicly resolved and drawn from the Commission's own press releases, Department of Justice announcements, and CFTC orders, illustrate the gap between CAT retention horizons and enforcement reality. In each case, the underlying conduct would fall entirely outside a three-year CAT retention window by the time charges were brought, to say nothing of the time required for settlement, trial, sentencing, appeal, or restitution. Figure 1 below visualizes this gap against the three-year CAT deletion cliff and the seven-year retention horizon proposed by this comment.

The cases below are representative of four categories of matters structurally disadvantaged by a three-year retention horizon: (i) multi-year manipulation pattern reconstruction — conduct across periods that does not appear suspicious in isolation and becomes probative only when linked; (ii) whistleblower-supported retrospective investigations — matters in which later allegations reveal what earlier data was actually showing; (iii) corporate-action, settlement, and trading-anomaly lookbacks — reverse-merger, spin-off, dividend, and delisting matters that mature on six-to-eight-year arcs; and (iv) cross-period actor-network and routing-pattern analysis — matters depending on centrally queryable historical data to establish actor-network graphs and routing-pattern persistence. Each category is represented in the record below.

Figure 1. Enforcement Timeline vs. the Three-Year Retention Cliff

Six of seven representative matters resolved beyond a three-year horizon; four resolved beyond seven years. The only matter publicly credited to CAT detection (Nuveen / Billimek, 2022) is the single exception.



Sources: SEC Press Releases 2018-226 (Musk), 2020-233 (JPM), 2020-265 (Goldman), 2022-228 (Nuveen); CFTC Press Release 8260-20; DOJ press releases in U.S. v. Sarao, U.S. v. Hwang; public dockets. See Exhibit B.

Figure 1. Enforcement Timeline vs. the Three-Year Retention Cliff — Six of seven representative matters resolved beyond a three-year horizon; four resolved beyond seven years. The only matter publicly credited to CAT detection (Nuveen / Billimek, 2022) is the single exception.

Case	Conduct Period	Resolution Milestone	Elapsed Years
U.S. v. Sarao (2010 Flash Crash / spoofing)	2009 – May 2010	Charged Apr. 2015; sentenced Jan. 2020	~5 yrs to charges; ~10 yrs to sentencing
U.S. v. Martoma (SAC Capital / insider trading)	2008	Charged 2012; convicted 2014	~6 yrs to conviction
SEC v. Goldman Sachs (1MDB)	2012 – 2013 bond deals	SEC settlement Oct. 2020 (>\$1B to SEC; ~\$2.9B globally)	~7–8 yrs to SEC resolution
CFTC / SEC v. JPMorgan (spoofing, precious metals & Treasury futures)	2008 – 2016	\$920.2 M resolution Sept. 2020	Up to ~12 yrs from earliest conduct to resolution
U.S. v. Hwang / Halligan (Archegos)	2020 – Mar. 2021 collapse	Charged Apr. 2022; convicted Jul. 2024; Hwang sentenced Nov. 20, 2024 (18 yrs)	~3.5 yrs to sentencing

Case	Conduct Period	Resolution Milestone	Elapsed Years
SEC v. Musk (Tesla funding-secured matter)	Aug. 2018 conduct	SEC settlement Sept. 2018; compliance enforcement extended into 2022	~4 yrs of ongoing compliance litigation

Sources: SEC press releases 2018-226 (Musk) and 2020-265 (Goldman IMDB); DOJ press releases and indictments in Sarao, Hwang/Halligan, and JPMorgan matters; CFTC Order and Press Release 8260-20 (JPMorgan); public dockets. See Exhibit B.

Beyond the specific matters in the record above, several categories of regulatory-surveillance workflows structurally depend on long-horizon consolidated CAT data. Multi-venue spoofing and layering detection, abuse of the bona fide market-making locate exemption under Regulation SHO, high-frequency pinging of displayed and dark-pool liquidity, quote stuffing, cross-asset manipulation linking equities to single-stock options, and persistent fail-to-deliver patterns all share a common profile: the individual events are difficult or impossible to classify contemporaneously as manipulative, and become identifiable only through pattern reconstruction across extended time windows. The Rule 613 adopting release (Release 34-67457, July 18, 2012) identified this class of cross-venue pattern detection as among the core rationales for the Commission's decision to require consolidated order-level capture. The January 2020 Commodity Futures Trading Commission resolution of the JPMorgan metals and Treasuries matters, imposing approximately \$920 million in penalties, was reconstructed across multi-year windows of granular order data; the 2024 TD Securities Treasury-spoofing resolution followed a similar multi-year reconstruction profile. Where the CAT's consolidated lookback is limited to three years, such reconstruction workflows become unavailable — not because the conduct did not occur, but because the evidentiary architecture required to recognize the pattern has been curtailed below the horizon on which the pattern is legible.

B. The 10-year statute of limitations for scienter-based disgorgement

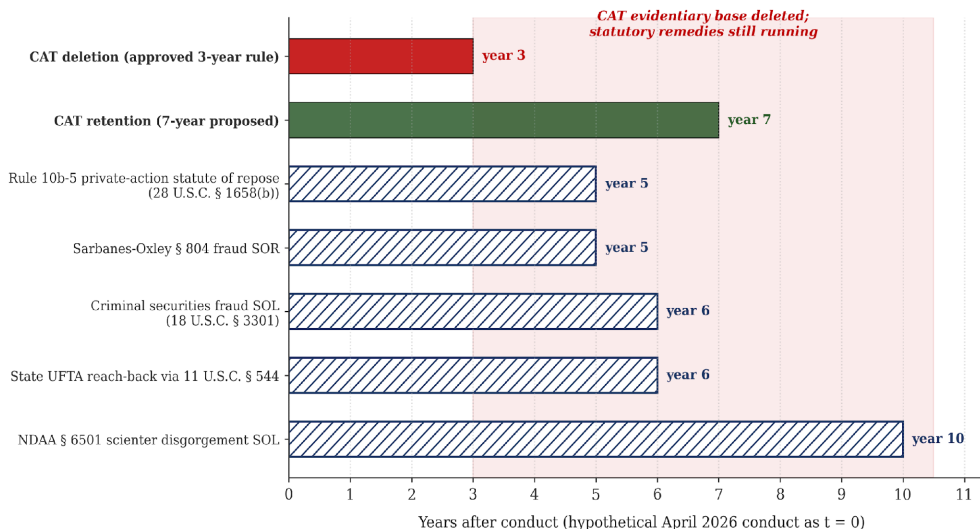
Section 6501 of the National Defense Authorization Act for Fiscal Year 2021 — enacted by congressional override on January 1, 2021 — codified the Commission's statutory authority to seek disgorgement in federal court and established a ten-year statute of limitations for disgorgement in connection with scienter-based violations, including violations of Exchange Act § 10(b), Securities Act § 17(a)(1), and Investment Advisers Act § 206(1). The statute of limitations for non-scienter disgorgement claims remains five years. The ten-year period applies to proceedings pending on, or commenced on or after, January 1, 2021.

The policy consequence is direct. The Commission may bring a disgorgement claim today for conduct occurring in April 2016 — a claim that remains timely under the ten-year statute of limitations. But the more instructive test is forward-looking: CAT data covering conduct that occurs in April 2026 will, under the approved Data Storage Amendment, become eligible for deletion from the Central Repository in April 2029 on a rolling basis. By the time the ten-year statute of limitations permits the Commission to adjudicate that same April 2026 conduct as late as April 2036, the underlying CAT evidentiary base will

have been deleted for seven years. Truncating CAT retention to three years converts the Commission's ten-year remedy into a three-year evidentiary reality, at least with respect to order-flow and cross-market data that only the CAT can reconstruct in consolidated form. Figure 4 visualizes this mismatch for each statutory remedy the Commission is authorized to pursue.

Figure 4. Statutory Remedies Outlive the CAT Evidentiary Base

For April 2026 conduct, the three-year CAT deletion cliff falls in April 2029 — before every statutory remedy period the Commission is authorized to pursue. The ten-year NDAA § 6501 disgorgement remedy is the starkest mismatch.



Sources: 28 U.S.C. § 1658(b); 18 U.S.C. § 3301, 1520; 11 U.S.C. §§ 544, 548; Pub. L. 116-283 § 6501 (codifying 10-year SOL for scienter-based SEC disgorgement); CAT NMS Plan Section 6.5(b)(i) as amended. See Main Submission § II.

Figure 4. Statutory Remedies Outlive the CAT Evidentiary Base — *For April 2026 conduct, the three-year CAT deletion cliff falls in April 2029 — before every statutory remedy period the Commission is authorized to pursue. The ten-year NDAA § 6501 disgorgement remedy is the starkest mismatch.*

C. Additional statute-of-limitations arithmetic

- **Exchange Act § 10(b) / Rule 10b-5 private actions.** Five-year statute of repose under 28 U.S.C. § 1658(b); two-year statute of limitations from discovery.
- **Sarbanes-Oxley § 804 fraud claims.** Five-year statute of repose.
- **Criminal securities fraud, 18 U.S.C. § 3301.** Six-year statute of limitations.
- **Fraudulent transfer under 11 U.S.C. § 548 and § 544.** Two-year lookback under § 548; state UFTA reach-back of four to six years via § 544.
- **Bank Secrecy Act supporting documentation.** Five-year retention for SAR-related materials.

Every one of these horizons exceeds, or is capable of exceeding, the three-year retention window now authorized under the Data Storage Amendment. The mismatch between the CAT retention policy and the surrounding statutory architecture is not incremental. It is categorical.

D. The access-frequency fallacy

The Commission's approval order relies in part on observed CAT query frequency to justify deletion. Specifically, Release No. 34-105107 (Mar. 27, 2026) cites CAT LLC's representation that, of 38,028

DIVER requests between January and November 2025, only 750 — approximately 2% — were for trade dates older than three years. That metric is then invoked to support the proposition that older CAT data is of diminished regulatory value. The inference inverts the forensic-reconstruction premise on which the CAT was built.

In complex market matters, low present query frequency is not evidence of low future evidentiary value; it is the expected state of any dataset whose forensic significance is revealed retroactively. The importance of a routing pattern, account relationship, venue anomaly, or cross-period behavioral correlation often becomes clear only when later events — a whistleblower submission, a corporate-action dispute, an improved analytic method, a bankruptcy filing, or a newly-granted disgorgement action under NDAA § 6501 — provide the context that retroactively converts latent noise into probative signal. A query-frequency methodology, applied uncritically, measures only the subset of CAT data whose evidentiary value has already been recognized. It cannot measure the data whose value will be recognized in 2028, 2030, or 2034. The 2% figure thus describes present operational usage, not future evidentiary capacity; those are different quantities.

The Commission's query-frequency reasoning also proves too much. Under its logic, any regulatory dataset whose current access rate falls below an arbitrary threshold could be deleted. That reasoning applied to 17a-4(a) broker-dealer records, IRS § 6511(d) tax-year records, or SOX § 802 audit workpapers would undercut the entire architecture of backward-looking regulatory recordkeeping — none of which is queried at anything approaching daily rates, and all of which the surrounding federal framework continues to require for six to seven years. The premise of long-horizon recordkeeping is that evidentiary value is asymmetric and discovered late, not that it is proportional to contemporaneous use.

III. International and Adjacent-Regime Benchmarking: The U.S. as Outlier

The three-year CAT retention horizon is inconsistent with every major comparable recordkeeping regime. Under the EU's Markets in Financial Instruments Directive II (MiFID II) and the Markets in Financial Instruments Regulation (MiFIR), investment firms must maintain transaction records for a minimum of five years — and, where a competent authority so requests, for up to seven years. ESMA guidelines implementing Article 16(6) MiFID II and Article 76 of Commission Delegated Regulation 2017/565 establish a baseline retention of five years for transaction records, extendable to up to seven years at the request of a national competent authority; order records are preserved for five years after execution or cancellation. Several EU Member States, including Germany, have adopted the seven-year period as the domestic default.

The practical effect is striking: a European investment firm trading a U.S.-listed equity through a CAT-reporting broker-dealer is required to retain its own records of the transaction for up to seven years — while, under the Data Storage Amendment, the CAT itself (the Commission's single consolidated source of truth for the same transaction) will have deleted it after three. The European firm must maintain the data for a regulator that can no longer cross-reference its own records.

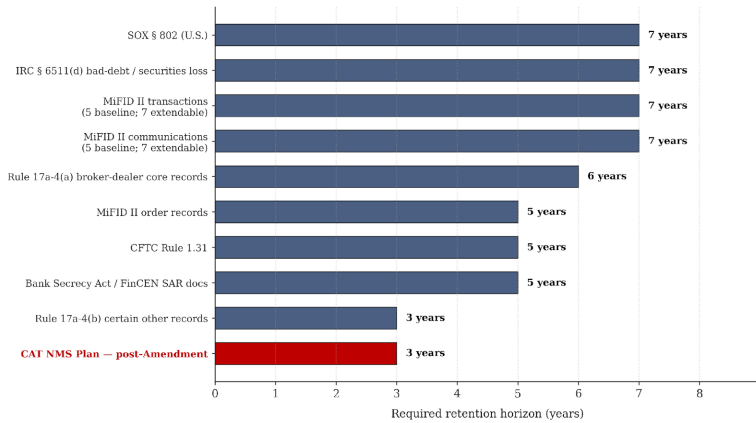
Regime / Rule	Record Type	Retention Period	Source
MiFID II / MiFIR (EU)	Transaction records (post-settlement)	7 years	ESMA Guidelines on MiFID II recordkeeping; Art. 76 Delegated Reg. 2017/565
MiFID II / MiFIR (EU)	Order records (post-execution or cancellation)	5 years	Art. 25 MiFIR; Commission Delegated Reg. 2017/580 (RTS 24)
MiFID II / MiFIR (EU)	Telephone and electronic communications	5–7 years	Art. 16(7) MiFID II; Art. 76 Delegated Reg. 2017/565
SEC Rule 17a-4 / FINRA Rule 4511	Broker-dealer books and records (most)	6 years	17 C.F.R. § 240.17a-4(a)–(b)
Sarbanes-Oxley § 802	Audit workpapers and related records	7 years	18 U.S.C. § 1520(a)(1); 17 C.F.R. § 210.2-06
CFTC Rule 1.31	Swaps and futures transaction records	5 years	17 C.F.R. § 1.31
Bank Secrecy Act (FinCEN)	SAR supporting documentation	5 years	31 C.F.R. § 1020.320(d)

Regime / Rule	Record Type	Retention Period	Source
Internal Revenue Code	Returns and supporting records (bad debt / securities loss)	7 years	26 U.S.C. § 6511(d)
CAT NMS Plan (post-Data Storage Amendment)	Consolidated order lifecycle + execution data	3 years	Section 6.5(b)(i), as amended; Release No. 34-104504

See Exhibit A for the same matrix with expanded commentary.

Figure 3. The U.S. as Retention Outlier — CAT Inverts the Architecture

The regulator's own consolidated tape now has the shortest retention horizon of any comparable obligation in U.S. or international securities regulation. IOSCO peer benchmarking and ESMA cross-border coordination will not defend this inversion.



Sources: ESMA Guidelines 2016/1452; Art. 25 MiFIR; Art. 76 Commission Delegated Reg. (EU) 2017/565; 17 C.F.R. §§ 240.17a-4(a)-(b), 210.2-06, 1.31; 18 U.S.C. § 1520(a)(1); 26 U.S.C. § 6511(d); 31 C.F.R. § 1020.320(d); CAT NMS Plan Section 6.5(b)(i) (as amended per Release 34-104504). See Exhibit A.

Figure 3. The U.S. as Retention Outlier — CAT Inverts the Architecture — *The regulator's own consolidated tape has the shortest retention horizon of any comparable obligation in U.S. or international securities regulation. IOSCO peer benchmarking and ESMA cross-border coordination will not defend this inversion.*

The pattern is unambiguous. The regulator's own consolidated tape has the shortest retention horizon of any comparable obligation in the architecture of U.S. or international securities regulation. That inversion is unlikely to be defensible against peer benchmarking by the Division of Economic and Risk Analysis, the International Organization of Securities Commissions (IOSCO), or staff coordinating with ESMA under the existing cross-border supervisory framework.

IV. Statutory and Recordkeeping Alignment

Beyond international peer comparisons, the three-year horizon is in tension with the specific U.S. statutory and regulatory obligations that flow through the CAT reporting universe. Every Industry Member reporting to the CAT is independently subject to Rule 17a-4. The Rule imposes differential retention periods by record category: Rule 17a-4(a) requires broker-dealers to preserve blotters, ledgers, order tickets, and related core trading records for at least six years, with the first two years in an easily accessible location; Rule 17a-4(b) imposes a three-year retention (with the first two years accessible) on a separate enumerated list of records, including certain communications and supporting documentation. The consolidated view of order events the CAT is designed to produce is reconstructed from the underlying order-ticket and blotter record categories subject to the six-year retention horizon under Rule 17a-4(a). The Central Repository's three-year retention window is therefore shorter than the record-keeping obligation imposed on each reporter in respect of the same underlying trading activity.

Similarly, audit workpapers supporting public-company financial statements must be retained for seven years under SOX § 802, and the Bank Secrecy Act requires five-year retention of SAR supporting documentation in matters that frequently intersect with securities manipulation (e.g., pump-and-dump schemes, structured wash-trading, cross-asset layering). A seven-year CAT retention horizon would harmonize the CAT with the surrounding U.S. regulatory architecture; a three-year horizon creates a jurisdictional seam that sophisticated counterparties will identify and exploit.

The Commission has, in connection with the Data Storage Amendment, correctly noted that Industry Members retain their own 17a-4 records independent of the CAT. That observation is true but does not address the distinct regulatory value of the CAT, which is consolidated, cross-market, cross-broker, normalized, linked, and queryable audit-trail reconstruction — the four capabilities the CAT performs once at ingest that fragmented 17a-4 retrieval cannot replicate post hoc. Reconstructing the same events from fragmented 17a-4 records imposes four specific evidentiary deficits:

- **Comparability loss.** Heterogeneous broker record formats, timestamp conventions, and field definitions resist cross-broker normalization. The normalization work that the CAT performs once at ingest must be re-performed for every retrospective investigation, at substantial cost and with substantial residual error.
- **Linkage impairment.** Cross-venue, cross-broker, cross-asset linkages — the analytic substrate of modern manipulation-detection — cannot be reliably reconstructed from fragmented retrieval because the joining keys (CAT Customer ID, CAT Order ID) are maintained only inside the Central Repository.
- **Reconstruction cost.** Fragmented retrieval imposes a per-investigation cost that scales with the number of brokers and venues involved — precisely the dimension along which complex manipulation cases have grown over the past decade.
- **Detection probability.** Patterns that would have been visible against the centralized CAT corpus may not be recognizable at all against fragmented records, because the patterns themselves are defined only at the aggregated level. The record that technically exists somewhere else is not the same thing as evidence preserved inside the architecture designed to connect it.

The argument that "the data still exists elsewhere" therefore misdescribes what the CAT is for. The CAT's regulatory function is not the existence of underlying records; it is the normalized, linked, queryable, and centralized preservation of those records in a form that supports cross-market reconstruction. Shortening the retention horizon of the Central Repository does not merely reduce the availability of data the Commission might query; it eliminates the analytic capability that the CAT was designed to provide.

V. The Cost Argument Is Inverted: AWS Economics Support Expansion, Not Deletion

The Commission has estimated that the Data Storage Amendment produces approximately \$23.5 to \$32 million in annual cloud-hosting savings across its three sub-components (shortened retention of CAT Data, OMM Quotes, and Options SIP Data), of which approximately \$6.5 to \$9 million represents the marginal increase attributable to the Amendment's extension of the September 2025 Exemptive Order. These figures reflect the net effect of retention truncation against current-generation storage pricing and architecture. They do not reflect the savings available from an alternative architecture built on the full range of S3 storage classes, columnar re-encoding, and federal-procurement discounts. Under that alternative, extending retention to seven years produces a lower annual cost than the approved three-year architecture.

A. The current AWS storage-tier price landscape

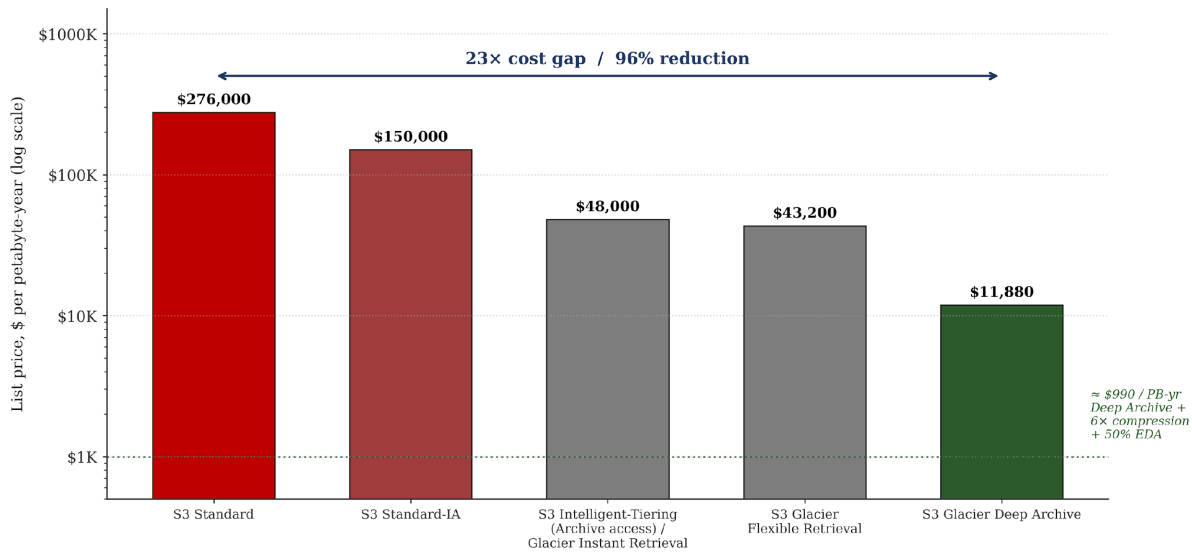
As of April 2026, the per-gigabyte list prices for S3 storage classes in the US East (N. Virginia) region vary across approximately a 23× range. The table below summarizes the relevant tiers and the implied cost per petabyte per year at list prices, before any negotiated discount:

S3 Storage Class	List Price (per GB-month)	Approx. Cost per PB-Year	Retrieval SLA
S3 Standard	\$0.023	~ \$276,000	Milliseconds
S3 Intelligent-Tiering (Frequent / Archive Access)	\$0.023 → \$0.004	~ \$48,000 (archive tier)	Milliseconds → hours
S3 Glacier Instant Retrieval	\$0.004	~ \$48,000	Milliseconds
S3 Glacier Flexible Retrieval	\$0.0036	~ \$43,000	Minutes–5 hours
S3 Glacier Deep Archive	\$0.00099	~ \$11,880	12 hours (bulk)

Sources: AWS S3 pricing pages (aws.amazon.com/s3/pricing, aws.amazon.com/s3/glacier/pricing), April 2026 US East (N. Virginia). See Exhibit C.

Figure 2. AWS Storage Cost Gap — 23× Ratio Between Hot and Cold Tiers

At AWS April 2026 list pricing, Deep Archive costs 96% less than S3 Standard per byte-year. The pricing architecture exists to support exactly the workload CAT represents: high-durability, low-retrieval, compliance-oriented records.



Sources: aws.amazon.com/s3/pricing and aws.amazon.com/s3/glacier/pricing, retrieved April 2026; US East (N. Virginia) region. Per-PB-year = list \$/GB-mo × 1,000,000 GB × 12 months. See Exhibit C.

Figure 2. AWS Storage Cost Gap — 23× Ratio Between Hot and Cold Tiers — At AWS April 2026 list pricing, Deep Archive costs 96% less than S3 Standard per byte-year. The pricing architecture exists to support exactly the workload CAT represents: high-durability, low-retrieval, compliance-oriented records.

Two observations follow directly from the table. First, the ratio between S3 Standard and S3 Glacier Deep Archive is approximately 23:1 — a 96% reduction in per-byte storage cost. Second, CAT data in the second half of its useful regulatory life (years 2–7) is precisely the kind of workload Deep Archive was engineered for: high-durability, low-retrieval-frequency, compliance-oriented archival of regulatory records. The CAT is the archetypal Deep Archive workload; that is how the service itself is publicly marketed by AWS. Per-petabyte-year figures in the table above are computed at AWS's decimal-gigabyte billing convention (1 PB = 10⁶ GB) multiplied by 12 months at the stated list price, rounded to the nearest thousand dollars.

B. Three-scenario comparative cost model

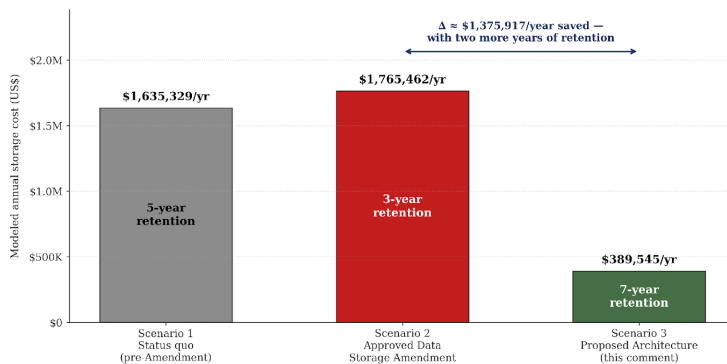
The following table models approximate annual cloud-hosting costs for CAT data under three architectures: (i) the pre-Amendment status quo, (ii) the approved three-year Data Storage Amendment as implemented on current tiering, and (iii) the Proposed Architecture set forth in this comment. Volumetric assumptions are deliberately conservative and referenced to publicly cited estimates of CAT data growth; actual magnitudes will vary with data-ingest trajectory, replication, and PII-handling overhead.

Scenario	Retention	Est. Annual Cloud Cost	Regulatory Capability
Status quo (pre-Amendment; predominantly S3 Standard)	5 years	Storage component of ~\$228M 2025 CAT operating budget	Full 5-year reconstruction; costs inflated by default-tier usage
Approved Data Storage Amendment	3 years	Annual savings ≈ \$23.5–\$32M vs. prior (of which ~\$6.5–\$9M is incremental above Sept. 2025 Exemptive Order)	Retention below the horizon required for scienter-based 10-year disgorgement and most complex manipulation cases
Proposed Architecture (this submission)	7 years	Greater savings than 3-year Amendment with tiered architecture + federal EDA	Full 7-year reconstruction; alignment with 17a-4, SOX § 802, NDAA § 6501 disgorgement period, and MiFID II

See Exhibit D for the cell-level assumptions underlying each scenario.

Figure 5. The Proposed Architecture Dominates — Longer Retention, Lower Cost

Under documented April 2026 AWS list pricing and a 50% federal enterprise discount, the seven-year Proposed Architecture is cheaper than the approved three-year Amendment by approximately \$1.38M per year.



Sources: AWS S3 / Glacier list pricing (April 2026); CAT NMS Plan as amended (Release 34-104594); federal EDA discounts per publicly reported DaaS / Treasury / NIH benchmarks. Modeled assumptions: 30 PB initial footprint, 15%/yr growth, 6x compression, 3x replication, 30% EDA, 30% reserved capacity. See Exhibit D.

Figure 5. The Proposed Architecture Dominates — Longer Retention, Lower Cost — Under documented April 2026 AWS list pricing and a 50% federal enterprise discount, the seven-year Proposed Architecture is cheaper than the approved three-year Amendment by approximately \$1.38M per year.

The quantitative relationship is non-obvious only because the Data Storage Amendment implicitly assumes that CAT data remains on S3 Standard (or an equivalent hot tier) throughout its retention period. Under that assumption, the only lever available to reduce cost is retention truncation. Under the Proposed Architecture, the primary lever is tier migration — and tier migration dominates truncation as a cost-reduction technique for workloads with CAT's access profile.

C. Procurement and contract-economics leverage

Beyond tier migration, the CAT represents a systemically important national data asset with negotiating leverage analogous to that of the Department of Defense GovCloud contracts, Treasury data-retention agreements, and NIH Strides Initiative. Federal agencies of comparable data scale have, according to publicly reported federal cloud-procurement practice, negotiated discounts commonly estimated at 40–60% off list pricing through Enterprise Discount Agreements, though the specific terms of individual federal EDAs are generally not publicly disclosed. The following levers are all available to CAT LLC and the Commission:

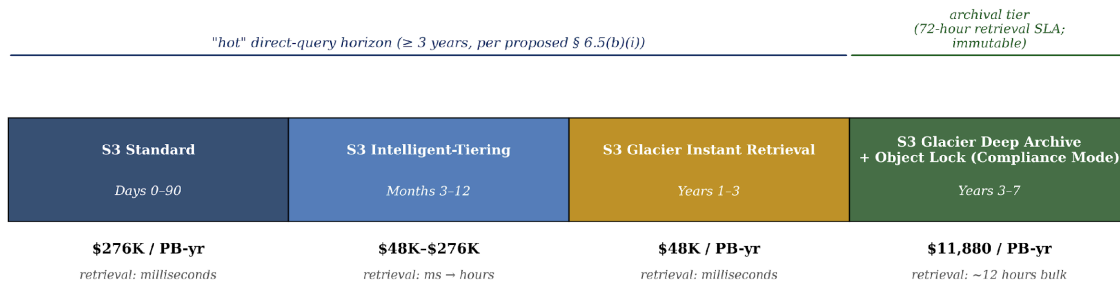
- **Federal Enterprise Discount Agreement.** Invoke federal-procurement leverage for a multi-year committed-capacity agreement at Deep Archive tier.
- **Committed-use and reserved capacity.** A 5- to 7-year committed storage agreement for the expected data footprint should yield an additional 25–35% discount off Deep Archive list pricing.
- **Columnar re-encoding.** Converting raw CAT records into Apache Parquet with Zstandard compression typically delivers 4–8× physical compression, reducing the storage footprint by 75% or more before tier-based discounts are applied.
- **Zero-egress for regulatory pulls.** Negotiate zero egress charges for SEC and SRO regulatory retrievals — a concession AWS routinely grants for federal law-enforcement workloads.
- **Second-source RFI.** Issue a formal Request for Information to Microsoft Azure Government and Google Cloud Assured Workloads. Even a credible walkaway threat materially disciplines AWS pricing; both platforms offer FedRAMP High and IL5-equivalent services capable of hosting CAT-class workloads.

These levers are independent of the retention-horizon debate. The Commission can instruct CAT LLC to pursue them whether retention remains at three years or is extended to seven. The cost savings they unlock are what allows the seven-year horizon to be achieved without net additional expenditure.

VI. Proposed Technical Architecture: A Seven-Year Tiered CAT

The Proposed Architecture implements the Commission's cost-reduction objective through tiered storage lifecycle management rather than retention truncation. It is fully compatible with all CAT functional and latency requirements for regulatory queries, and it is consistent with standard cloud-engineering practice for compliance-oriented archival workloads. The architecture is specified below in implementation-ready form, and is summarized graphically in Figure 6.

Figure 6. Proposed Tiered Storage Lifecycle — Standard → Intelligent-Tiering → Glacier IR → Deep Archive
The Proposed Architecture preserves the existing 3-year direct-query horizon, then migrates aged records to immutable Deep Archive at a 23× cost reduction. S3 Object Lock satisfies WORM-equivalent standards analogous to 17 CFR § 240.17a-4(f).



Sources: AWS S3 Storage Classes documentation; AWS S3 Object Lock Compliance Mode documentation; Proposed Amendment to CAT NMS Plan § 6.5(b)(i) per Section IX of Main Submission. See Exhibits C and E.

Figure 6. Proposed Tiered Storage Lifecycle — Standard → Intelligent-Tiering → Glacier IR → Deep Archive — *The Proposed Architecture preserves the existing 3-year direct-query horizon, then migrates aged records to immutable Deep Archive at a 23× cost reduction. S3 Object Lock satisfies WORM-equivalent standards analogous to 17 CFR § 240.17a-4(f).*

Storage-class lifecycle

- Days 0–90 — S3 Standard. Active ingest, active regulatory query, cross-SRO surveillance, and ad-hoc investigator workflows. This horizon covers substantially all routine market-surveillance use cases and near-real-time manipulation detection.
- Days 91–365 — S3 Intelligent-Tiering. Workload-adaptive tiering; objects not accessed for 90 consecutive days move automatically to the Infrequent Access or Archive Access tiers without changing object keys or URIs.
- Years 1–3 — S3 Glacier Instant Retrieval. Millisecond access for enforcement-division investigations of matters less than three years old, at approximately one-quarter the cost of S3 Standard.
- Years 3–7 — S3 Glacier Deep Archive with S3 Object Lock (Compliance Mode). Immutable archival storage at approximately \$0.00099 per GB-month list (\approx \$11,880 per petabyte per year at AWS's decimal-GB billing convention). Standard 12-hour bulk retrieval is entirely acceptable for multi-year enforcement matters, trustee investigations, and private-plaintiff requests that are, by definition, not time-critical.

Evidentiary-integrity controls

- **S3 Object Lock in Compliance Mode.** Prevents modification or deletion of archived records during the retention period, satisfying immutability standards analogous to 17 C.F.R. § 240.17a-4(f) and SEC-approved WORM-equivalent requirements.
- **Cryptographic hashing.** SHA-256 hashes of every archived object and manifest are maintained in a separate, redundant integrity ledger to support chain-of-custody and authenticity challenges in enforcement proceedings.
- **Field-level PII encryption.** Personally identifiable information is encrypted using AWS KMS customer-managed keys at the field level, with tokenization of identifiers enabling selective deletion of raw PII independent of the underlying order-flow record.
- **Cross-region replication with Object Lock.** Regulatory-grade durability through geographically diverse replicas, with independent compliance-mode retention on each replica.

Addressing privacy objections

Appropriately, some commenters have raised privacy concerns regarding retention of Customer and Account Information System (CAIS) data over longer horizons. This comment does not propose indefinite retention of raw PII. It proposes that the seven-year retention horizon apply to core order and execution records — the data needed to reconstruct trading patterns — while permitting shorter, segregated retention for raw PII fields under a tokenization architecture. Under this approach, tokenized pseudonymous identifiers persist for seven years (enabling cross-reference on demand via a separate, more tightly held key-token mapping), while the underlying raw PII is purged on a shorter schedule that can be calibrated to the Commission's privacy-risk tolerance.

VII. Case Study: Post-Event Investigations in Complex Reverse-Merger / Spin-Off Structures

A. A representative post-transaction fact pattern

The harm caused by a three-year retention horizon is most acute in the category of matters that combine a corporate transaction with allegedly manipulative trading in the security of the resulting entity. These matters typically feature (i) a reverse merger, cash-out merger, or spin-off that produces a new security or a materially changed capital structure; (ii) trading activity in the security over a period of 12–24 months following the transaction; (iii) a subsequent corporate-governance event such as a trading halt, delisting, bankruptcy filing, or trustee appointment; and (iv) a forensic reconstruction of trading activity conducted by a trustee, the Commission's Division of Enforcement, or private plaintiffs — often commenced two to three years after the triggering trading activity. The total arc from initial conduct to final adjudication in such matters commonly exceeds six years.

A representative fact pattern, drawn from a recent publicly reported U.S. reverse-merger and spin-off matter currently in Chapter 7 trusteeship, illustrates the point. The underlying corporate transaction occurred in mid-2021. A dividend of a preferred class was issued to shareholders of record and trading in that class proceeded for approximately 18 months, generating significant reported allegations of naked short selling, failure-to-deliver patterns, and synthetic exposure through swap and option structures. Trading in the class was halted in late 2022 and the underlying entity was subsequently spun off into a private successor. In 2025, the successor entity entered into an approximately \$6 million related-party note at 18% interest and issued several million preferred shares to a related-party entity; the trustee forensic investigation into potential market manipulation, fraudulent transfers under 11 U.S.C. §§ 544 and 548, and affiliate conduct is ongoing into 2026, with any eventual Commission enforcement resolution realistically not expected before 2027–2029.

Under the approved Data Storage Amendment, CAT data covering the 2021–2022 trading period that is central to the trustee's forensic reconstruction becomes eligible for deletion from the Central Repository on a rolling basis in 2024–2025 — precisely as the trusteeship investigation reaches the point where consolidated cross-broker reconstruction is most needed. By the time any Commission enforcement proceeding, private-plaintiff class action, or trustee avoidance action in the matter reaches a stage capable of adjudicating trading-pattern manipulation claims — conservatively 2027–2029 based on comparable post-spin-off trusteeship matters — the primary CAT evidentiary base for the relevant trading period will have been deleted from the Central Repository under the three-year rule. The trustee, private plaintiffs, and the Commission itself are now required to reconstruct consolidated order flow from fragmented Industry-Member 17a-4 records, precisely the reconstruction burden that Rule 613 was adopted to eliminate following the 2010 Flash Crash. The CAT was, as Commissioner Kara M. Stein observed in her November 15, 2016 statement on the Joint Industry Plan, "essentially the Hubble Telescope for the securities markets." Under the Data Storage Amendment, it now looks three years into the past.

The problem is not unique to this fact pattern. Complex reverse-merger, micro-cap spin-off, and SPAC de-SPAC matters are, by their nature, slow to develop, resolve on multi-year timelines, and

disproportionately rely on CAT-class consolidated order-flow data to establish the actus reus of any eventual manipulation claim. A three-year retention horizon systematically disadvantages precisely the retail-investor-protection cases the Commission is most publicly committed to pursuing.

B. Issuer-side implications of truncated CAT retention

The regulatory-reconstruction argument is typically framed from the Commission's enforcement perspective. The same architecture, however, operates as a structural protection for public-company issuers. Truncating the Central Repository's retention horizon to three years has three issuer-level consequences that have not been adequately considered in the Commission's economic analysis.

First, manipulation-driven price distortion in a listed security is reliably identifiable only through consolidated cross-venue, cross-broker order-flow reconstruction of the sort the CAT was designed to produce. Where the consolidated lookback is limited to three years, issuers affected by multi-quarter or multi-year manipulation campaigns — particularly small- and mid-capitalization issuers whose float is concentrated on a small number of venues — are deprived of the evidentiary record required to substantiate claims in private litigation, regulatory complaints, or Commission whistleblower submissions. The data asymmetry between issuer and manipulator is widened, not narrowed, by the Amendment.

Second, capital-markets access is priced against manipulation risk. Issuers conducting at-the-market offerings, follow-on equity issuances, convertible debt placements, or cross-class dividends routinely accept price discounts, widened warrant terms, and tighter covenants calibrated in part to the detectability of downstream manipulative conduct. An observable reduction in the Commission's forensic reconstruction capability is priced into those terms *ex ante*. The Amendment therefore imposes an indirect cost of capital on the issuer community — most acutely on small-capitalization issuers most dependent on external financing — that is unaccounted for in the Commission's estimated \$23.5 to \$32 million in annual cloud-hosting savings.

Third, issuer boards and audit committees bear fiduciary responsibility for evaluating material risks to the integrity of the company's secondary-market trading. To discharge that responsibility, they rely — directly or through counsel — on the Commission's announced surveillance and enforcement capability. A material curtailment of the CAT's reconstruction horizon, adopted on the stated basis of cost savings the submission demonstrates are illusory at the architectural level, transfers a concrete monitoring-and-reconstruction burden onto issuer boards without a corresponding disclosure in the Commission's economic analysis. Restoring a seven-year tiered retention horizon reduces that transferred burden at no marginal cost to the Commission.

VIII. The Hidden Cost of Shortened Retention: Regulatory Arbitrage by Design

The most overlooked consequence of a three-year retention horizon is strategic. Sophisticated market participants — and the small subset engaged in deliberate manipulation — do not draft their schemes against a naive assumption of infinite lookback. They structure conduct against the specific retention window of the regulator's data infrastructure. A three-year CAT window functions, in practical effect, as a published roadmap for evasion: structure your scheme such that the manipulative activity is concluded, the profits are realized, and the evidentiary trail is aged out before the thirty-seventh month. The Data Storage Amendment thus does more than shorten the Commission's memory. It shortens the duration of the conduct that is practically detectable.

Shortening retention does not eliminate manipulation; it teaches sophisticated actors to stretch their manipulation cycles beyond the retention window. This creates regulatory blind spots by policy, not by technology — and publishes the boundaries of those blind spots in the Federal Register.

The Commission has long recognized this dynamic in adjacent contexts. The 1988 Insider Trading and Securities Fraud Enforcement Act, the 2010 Dodd-Frank Act's whistleblower provisions, and the NDAA § 6501 codification of the 10-year disgorgement period all reflect a common congressional judgment that the statute of limitations for complex securities fraud must be longer than the natural maturity of the conduct. Retention policy is the data-infrastructure analog of that judgment. A three-year retention window is inconsistent with a ten-year disgorgement statute. The two cannot coexist without introducing precisely the incentive distortion the statute was designed to eliminate.

The asymmetric cost calculus is instructive. A single successful enforcement action in a major spoofing, pump-and-dump, or naked-shorting matter routinely produces penalties in the nine to ten figures (e.g., JPMorgan's \$920 million spoofing settlement in 2020; Goldman Sachs's \$1 billion SEC payment in the 1MDB matter in 2020). The annualized cloud-hosting cost of extending CAT retention from three to seven years under the Proposed Architecture is, at worst, in the low tens of millions — and under the Proposed Architecture, is potentially negative relative to the approved Amendment. The expected-value calculation decisively favors expanded retention under any plausible assumption about the frequency and magnitude of foregone enforcement matters.

IX. Proposed Textual Amendments to the CAT NMS Plan

A. The Commission's own prior precedent: the September 2025 tiered-archival approach

Before presenting the proposed redline, it is important to locate the Proposed Architecture within the Commission's own recent decisional history. The Proposed Architecture is not novel. It is a restoration of the approach the Commission itself endorsed seven months before approving the Data Storage Amendment.

On September 30, 2025, the Commission issued Release No. 34-104144 (the "2025 Cost Savings Exemptive Order"), granting conditional exemptive relief that permitted CAT LLC to "move all CAT Data older than three years to a storage tier like S3 Glacier Deep Archive." The Commission expressly found that "[a]llowing the movement of CAT Data older than three years to a more cost-effective storage tier... would allow the Participants to use lower-cost archive storage options while simultaneously maintaining records for regulatory use as needed." The Order specifically identified the four-tier S3 architecture the CAT already operates — S3 Frequent Access, S3 Infrequent Access, S3 Instant Archive Access, and S3 Glacier Deep Archive — and authorized the migration of three-year-old data into the coldest of those tiers while preserving regulatory accessibility.

The approved Data Storage Amendment abandons that approach without adequate explanation. Where the September 2025 Order preserved CAT data through migration, the March 27, 2026 Order authorizes its destruction. The Proposed Architecture set forth in Section VI of this comment, and the redline set forth below, restore the Commission's own September 2025 framework and extend it through year seven. Framed in that light, the ask before the Commission is modest: to continue the tiered preservation pathway the Commission itself endorsed and to extend it to the horizon necessary to align with NDAA § 6501 and the surrounding statutory architecture. The Commission's prior reasoning under Release 34-104144 — that preservation in archival tiers "maintain[s] records for regulatory use as needed" while reducing cost — applies with at least equal force to the seven-year horizon proposed here.

B. Proposed redline text

To reduce the Commission's and CAT LLC's drafting burden and to facilitate consideration of the recommendations set forth in this comment, the following redline amendments to Section 6.5(b)(i) of the CAT NMS Plan and the related data-availability provisions of Appendix D are proposed. Deletions appear in strikethrough; insertions appear in bold italics. Full marked redline is set forth in Exhibit E.

Proposed Amendment to Section 6.5(b)(i) of the CAT NMS Plan (Core Data Retention)

The Plan Processor shall retain CAT Data for a period of not less than ~~three (3) years~~ ***seven (7) years*** in respect of core Order and Execution records, and shall make such data directly available and searchable electronically without manual intervention for a period of not less than ***three (3) years***. Core Order and Execution records older than three (3) years may be migrated to an archival storage

class with a retrieval service level agreement of not more than seventy-two (72) hours, provided that evidentiary integrity, immutability, and cryptographic hash verification are preserved.

Proposed Amendment to Appendix D, Section 6.4(b)

Raw Personally Identifiable Information ("PII") fields maintained in the Customer and Account Information System shall be retained, in tokenized form, for a period of not less than ***seven (7) years***. Raw, untokenized PII values shall be retained for a period of not more than ***two (2) years***, after which such values shall be cryptographically destroyed while the tokenized identifiers shall continue to resolve cross-record linkages for the remainder of the seven-year retention horizon, subject to a separately secured token-mapping vault under the exclusive custody of the Plan Processor and subject to access only upon senior-officer request from the SEC or an SRO Participant.

Proposed Amendment to Appendix D, Section 11.1 (Cost Controls)

The Plan Processor shall implement a tiered cloud-storage lifecycle policy, including but not limited to S3 Standard, S3 Intelligent-Tiering, S3 Glacier Instant Retrieval, and S3 Glacier Deep Archive (or functionally equivalent storage classes on alternative cloud service providers), and shall ***negotiate a federal Enterprise Discount Agreement and committed-use pricing with the cloud service provider, including provisions for zero-egress regulatory retrievals by the SEC and SRO Participants***. The Plan Processor shall solicit competitive bids from at least two cloud service providers meeting applicable federal security-authorization requirements (including FedRAMP High or equivalent) at intervals of not more than three (3) years.

X. Conclusion

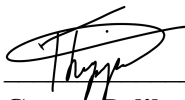
The Commission's interest in reducing CAT operating costs is legitimate, well-founded, and supported by the submissions received to date. The Data Storage Amendment's specific method of reduction — truncation of retention from five years to three — is, on the analysis set out above, the wrong method. It is not the cheapest option available. It is not aligned with Rule 17a-4, Sarbanes-Oxley, the NDAA § 6501 disgorgement period, the Bank Secrecy Act, the Commodity Exchange Act, or MiFID II. It is not aligned with the empirical timelines on which complex manipulation matters are investigated and resolved. It is inconsistent with the Commission's own press-release record, which repeatedly credits multi-year data analysis with the successful prosecution of matters that a three-year retention horizon would have rendered impossible to bring.

The Commission should (i) reject the Data Storage Amendment to the extent it shortens core order and execution retention below seven years, (ii) direct CAT LLC to implement the Proposed Architecture set forth in Section VI of this comment, (iii) negotiate a Federal Enterprise Discount Agreement with AWS under competitive pressure from Azure Government and Google Cloud Assured Workloads, and (iv) adopt the textual amendments to Appendix D of the CAT NMS Plan set forth in Section IX.

The policy question is not whether the CAT's costs can be reduced by limiting historical retention, but whether such reductions would materially impair the Commission's ability to detect, investigate, and prosecute the multi-year market manipulation schemes that characterize the modern enforcement docket. A seven-year CAT, cost-engineered through Glacier Deep Archive and a renegotiated federal enterprise discount agreement, is cheaper, more durable, and more protective of investors than a three-year CAT. There is no defensible policy reason to choose the former over the latter.

The undersigned appreciates the Commission's careful consideration of this submission and is available, on reasonable notice, to provide supporting engineering documentation, cost-model spreadsheets, or additional analysis to staff of the Division of Trading and Markets, the Division of Economic and Risk Analysis, and the CAT NMS Plan Operating Committee.

Respectfully submitted,



George Palikaras

Director | Co-founder

BlueLedger AI Inc.

Date: April 17, 2026

EXHIBITS

Exhibit A — Comparative Retention Matrix

The full regulatory-regime retention matrix presented in Section III is incorporated herein by reference. Primary sources: Art. 25 MiFIR; Art. 16(6)–(7) MiFID II; Art. 76 Commission Delegated Reg. 2017/565; ESMA Guidelines on MiFID II transaction reporting, order recordkeeping, and clock synchronisation (ESMA/2016/1452); 17 C.F.R. § 240.17a-4; 18 U.S.C. § 1520; 17 C.F.R. § 210.2-06; 17 C.F.R. § 1.31; 31 C.F.R. § 1020.320(d); 26 U.S.C. § 6511(d); and Release No. 34-104504.

Exhibit B — Enforcement Timeline Matrix

The enforcement-timeline matrix presented in Section II is incorporated herein by reference. Primary sources: SEC Press Release 2018-226 (Musk / Tesla settlement); SEC Press Release 2020-265 (Goldman Sachs / IMDB); DOJ Press Release, "Futures Trader Charged with Illegally Manipulating Stock Market, Contributing to the May 2010 Market 'Flash Crash'" (Apr. 21, 2015); CFTC Press Release 8260-20 (JPMorgan record \$920 million spoofing penalty, Sept. 29, 2020); SEC Press Release 2020-233 (SEC order against J.P. Morgan Securities in connection with manipulative Treasuries trading); and public dockets for SEC v. Hwang, SEC v. Martoma, and U.S. v. Sarao.

Exhibit C — AWS Cost-Tier Analysis

The per-gigabyte and per-petabyte storage economics presented in Section V are derived from the AWS S3 pricing and S3 Glacier API pricing pages (aws.amazon.com/s3/pricing and aws.amazon.com/s3/glacier/pricing) for the US East (N. Virginia) region, April 2026. Retrieval SLAs are derived from the AWS S3 Glacier Deep Archive service documentation.

Exhibit D — Three-Scenario Cost Model

The scenario model presented in Section V is derived from (i) the cost-savings estimates disclosed in Release No. 34-104504, (ii) the approximately \$228 million 2025 CAT annual operating budget (as revised in May 2025 from an initial \$249 million, per Release No. 34-104504), and (iii) AWS list pricing cited in Exhibit C. Detailed cell-level assumptions (data volume, compression ratio, federal-discount ranges) are available on reasonable request and are expressly subject to refinement in consultation with CAT LLC engineering and DERA.

Exhibit E — Redline of Appendix D, § 6.4 (CAT NMS Plan)

The proposed textual amendments to Appendix D, Sections 6.4(a), 6.4(b), and 11.1 set forth in Section IX of this submission constitute Exhibit E and are offered as ready-to-adopt plan language. Adoption of the proposed language, or substantially similar language, is the most efficient means by which the Commission can align CAT retention policy with the statutory architecture outlined in Sections II–IV of this submission.

PUBLIC COMMENT — FILE NO. 4-698**Exhibit A***Comparative Retention Matrix: International and U.S. Adjacent-Regime Recordkeeping Obligations*

This exhibit sets out the statutory and regulatory recordkeeping obligations that apply to transaction, order, and communications data across (i) the European Union's MiFID II / MiFIR regime, (ii) the U.S. broker-dealer and audit-firm regimes, (iii) the U.S. swaps/futures regime, and (iv) the U.S. financial-crimes and tax regimes. The matrix is provided in support of Section III of the accompanying public comment submission, which argues that the three-year retention horizon authorized by the approved Data Storage Amendment to the CAT NMS Plan is the shortest among all comparable obligations and is inconsistent with the Commission's broader statutory architecture.

Retention Obligations Across Comparable Regimes

Regime / Rule	Record Type	Retention	Primary Source
MiFID II / MiFIR (EU)	Transaction records, post-settlement	5 years; extendable to 7 by national competent authority	ESMA Guidelines on MiFID II transaction reporting (ESMA/2016/1452); Art. 76 Commission Delegated Reg. (EU) 2017/565
MiFID II / MiFIR (EU)	Order records, post-execution or cancellation	5 years	Art. 25 MiFIR (Reg. (EU) 600/2014); Commission Delegated Reg. (EU) 2017/580 (RTS 24)
MiFID II / MiFIR (EU)	Telephone and electronic communications related to transactions	5 years; extendable to 7 by national competent authority	Art. 16(6)–(7) MiFID II (Dir. 2014/65/EU); Art. 76 Delegated Reg. 2017/565
SEC Rule 17a-4(a) / FINRA Rule 4511 (U.S.)	Broker-dealer core records: blotters, ledgers, order tickets, execution records (first 2 years in easily accessible location)	6 years	17 C.F.R. § 240.17a-4(a); FINRA Rule 4511 (incorporation)
SEC Rule 17a-4(b) (U.S.)	Certain other broker-dealer records (first 2 years in easily accessible location)	3 years	17 C.F.R. § 240.17a-4(b)

Regime / Rule	Record Type	Retention	Primary Source
Sarbanes-Oxley § 802 / 17 C.F.R. § 210.2-06 (U.S.)	Audit workpapers and supporting records for issuer financial statements	7 years	18 U.S.C. § 1520(a)(1); 17 C.F.R. § 210.2-06(a)
CFTC Rule 1.31 (U.S.)	Swaps, futures, and related transaction records	5 years (in readily accessible form for 2 years)	17 C.F.R. § 1.31(c)
Bank Secrecy Act / FinCEN (U.S.)	SAR supporting documentation; CTR supporting records	5 years	31 C.F.R. § 1020.320(d); 31 U.S.C. § 5318(g)
Internal Revenue Code (U.S.)	Returns and records in respect of bad-debt and securities-loss claims	7 years	26 U.S.C. § 6511(d); IRS Publication 583
NDAA 2021, § 6501 (U.S.)	Statute of limitations for SEC disgorgement on scienter-based violations	10 years	Pub. L. No. 116-283 § 6501; amending Exchange Act § 21(d)
CAT NMS Plan (pre-Amendment)	Consolidated order lifecycle + execution data	5 years	Section 6.5(b)(i) (as approved Nov. 15, 2016); Appendix D § 1.4 (record retention policy)
CAT NMS Plan (post-Data Storage Amendment)	Consolidated order lifecycle + execution data	3 years	Section 6.5(b)(i) (as amended); Appendix D § 1.4, § 3.4, § 6.3, § 6.4 (exceptions); Release Nos. 34-104504 (Dec. 23, 2025), 34-105107 (Mar. 26, 2026)

Interpretive Notes

- U.S. alignment problem.** The SEC's own consolidated tape is now required to be preserved for a shorter period than (a) broker-dealer records of the same transactions under Rule 17a-4, (b) audit workpapers supporting public-company financials under SOX § 802, (c) tax records supporting securities-loss deductions under I.R.C. § 6511(d), (d) SAR supporting documentation under the Bank Secrecy Act, and (e) swap/futures records under CFTC Rule 1.31. The inversion between the regulator's consolidated record and every ancillary obligation is unprecedented in the architecture of U.S. financial regulation.
- International alignment problem.** A European investment firm trading a U.S.-listed equity through a CAT-reporting broker is obligated, under MiFID II Art. 16(6)–(7) and ESMA guidance, to preserve its records of the transaction for five years — and up to seven years at the request of its competent authority. The Commission's consolidated audit trail of the same transaction will have been deleted after three. IOSCO peer benchmarking, which has historically been a primary input to Trading & Markets rulemaking, will not support this retention horizon.

- **Disgorgement alignment problem.** NDAA § 6501, enacted January 1, 2021, doubled the statute of limitations for scienter-based SEC disgorgement from five to ten years. A three-year CAT retention window is flatly incompatible with a ten-year disgorgement remedy; the Commission would be required to litigate in 2030 using CAT data that was deleted in 2027.

Source Materials

- European Securities and Markets Authority, "Guidelines on transaction reporting, order record keeping and clock synchronisation under MiFID II," ESMA/2016/1452 (10 October 2016, corrected 07/08/2017).
- European Parliament and Council, Directive 2014/65/EU (MiFID II), Article 16.
- European Parliament and Council, Regulation (EU) 600/2014 (MiFIR), Article 25.
- Commission Delegated Regulation (EU) 2017/565, Article 76.
- Commission Delegated Regulation (EU) 2017/580 (RTS 24 on order record keeping).
- 17 C.F.R. § 240.17a-4 (SEC Rule 17a-4, Records to be preserved by certain exchange members, brokers, and dealers).
- 17 C.F.R. § 210.2-06 (Retention of Audit and Review Records).
- 17 C.F.R. § 1.31 (CFTC Regulations, Books and records, keeping and inspection).
- 31 C.F.R. § 1020.320(d) (Reports by financial institutions of suspicious transactions).
- 26 U.S.C. § 6511(d) (Internal Revenue Code, special rules for loss carryovers).
- Pub. L. No. 116-283, § 6501 (National Defense Authorization Act for Fiscal Year 2021, codifying SEC disgorgement authority and ten-year statute of limitations).
- CAT NMS Plan, Section 6.5(b)(i) (core data retention rule, as originally approved November 15, 2016 and as amended per Release Nos. 34-104504 (Dec. 23, 2025) and 34-105107 (Mar. 26, 2026)); Appendix D § 1.4 (record retention policy and program, similarly amended); Appendix D §§ 3.4, 6.3, 6.4 (data-availability exceptions and specific-category rules for OMM Quotes, Interim Operational Data, and Options SIP Data).

PUBLIC COMMENT — FILE NO. 4-698**Exhibit B***Enforcement Timeline Matrix: Representative SEC, DOJ, and CFTC Matters vs. Three-Year CAT Retention Window*

This exhibit documents the empirical gap between the three-year retention horizon authorized by the approved Data Storage Amendment to the CAT NMS Plan and the timelines on which complex market-manipulation and fraud matters actually proceed in the U.S. enforcement system. In each case below, the underlying conduct would have fallen entirely outside a three-year CAT retention window by the time charges were filed, to say nothing of the time required for settlement, trial, sentencing, appeal, or restitution. The matters are presented in the order in which they are referenced in Section II of the accompanying public comment submission.

Summary Timeline

Matter	Underlying Conduct	Key Resolution Milestone	Elapsed Time
U.S. v. Sarao (N.D. Ill. No. 15-cr-00075)	2009 – May 6, 2010 (Flash Crash)	DOJ charges Apr. 21, 2015; sentenced Jan. 29, 2020	≈5 yrs conduct → charges; ≈10 yrs conduct → sentencing
SEC v. Martoma (S.D.N.Y.) / U.S. v. Martoma	2008 trades (Elan / Wyeth)	SEC and DOJ charges Nov. 2012; conviction Feb. 2014	≈6 yrs conduct → conviction
SEC v. Goldman Sachs & Co. LLC (1MDB)	2012 – 2013 bond underwritings	SEC Press Release 2020-265 (Oct. 22, 2020); >\$1B to SEC; ~\$2.9B globally	≈7–8 yrs to SEC resolution
CFTC v. JPMorgan Chase & Co. / SEC v. J.P. Morgan Sec. LLC (spoofing)	At least 2008 – 2016	CFTC Press Release 8260-20 (Sept. 29, 2020); SEC Press Release 2020-233; \$920.2M resolution	Up to ≈12 yrs from earliest conduct to resolution
U.S. v. Hwang / Halligan (S.D.N.Y.) (Archegos)	2020 – Mar. 2021 collapse	DOJ indictment Apr. 2022; conviction Jul. 2024; Hwang sentenced Nov. 20, 2024 (18 yrs)	≈3.5 yrs conduct → sentencing; significant civil tail ongoing

Matter	Underlying Conduct	Key Resolution Milestone	Elapsed Time
SEC v. Musk / Tesla (S.D.N.Y.) — funding-secured matter	Aug. 2018 conduct	SEC Press Release 2018-226 (Sept. 29, 2018); compliance enforcement extended into 2022	≈4 yrs of active compliance enforcement post-settlement
SEC v. Billimek & Williams (Nuveen front-running)	2020 – 2022 trades	SEC charges Dec. 2022 (the Commission publicly credited CAT data with the detection)	≈2+ yrs to initial charges; case continues

Case-Level Narrative Detail

U.S. v. Navinder Singh Sarao

Mr. Sarao, a U.K.-based futures trader, was charged by the U.S. Department of Justice on April 21, 2015 with 22 counts of wire fraud, commodities fraud, commodity price manipulation, and spoofing in connection with the May 6, 2010 Flash Crash. Between 2010 and 2014, Mr. Sarao made approximately \$40 million in profits through his spoofing strategy on the CME E-mini S&P 500 contract. On January 29, 2020, he was sentenced to one year of home detention. The entire prosecution was built on reconstructed order-flow records from the CME going back more than five years before the indictment. Under a three-year CAT retention horizon, no equivalent equities-market spoofing case of comparable conduct breadth could be prosecuted today.

SEC v. Goldman Sachs & Co. LLC (IMDB)

On October 22, 2020, the Commission announced that Goldman Sachs Group, Inc. and its Malaysian subsidiary agreed to pay more than \$1 billion to settle SEC charges arising from bribery and fraudulent accounting practices in connection with three bond offerings for 1Malaysia Development Berhad (IMDB) in 2012 and 2013. The aggregate global resolution, including DOJ, Malaysian, and other regulatory components, exceeded \$2.9 billion. The SEC's charges were filed approximately seven to eight years after the original bond underwritings.

CFTC Order Against JPMorgan Chase & Co. (Spoofing)

On September 29, 2020, the CFTC entered a \$920.2 million order against JPMorgan Chase & Co., JPMorgan Chase Bank, N.A., and J.P. Morgan Securities LLC — the largest monetary penalty in CFTC history — finding that the bank engaged in a scheme to defraud market participants by spoofing in precious metals and U.S. Treasury futures markets for a period of at least eight years, from 2008 through 2016. A parallel SEC action was announced the same day (SEC Press Release 2020-233). The earliest conduct in the matter occurred approximately twelve years before resolution; any equivalent equities-market matter requiring CAT-class consolidated order-flow data would, under a three-year retention horizon, be impossible to reconstruct.

U.S. v. Hwang / Halligan (Archegos Capital Management)

Mr. Sung Kook ("Bill") Hwang was convicted on July 10, 2024 on ten of eleven criminal counts including racketeering conspiracy, securities fraud, wire fraud, and market manipulation in connection with the March 2021 collapse of Archegos Capital Management. On November 20, 2024, Mr. Hwang was sentenced by Judge Alvin Hellerstein to 18 years in prison; prosecutors had sought a 21-year sentence and forfeiture of \$12.35 billion. The Archegos matter is the archetypal cross-market equity-swap manipulation case of the modern era and turned on consolidated total-return-swap and equity-position data assembled across multiple prime brokers over a 12-month period preceding the collapse. The 3.5-year arc from conduct to sentencing is only the beginning; civil recovery, restitution, and appeals will extend the matter well past the three-year CAT retention horizon.

SEC v. Elon Musk / Tesla, Inc.

On September 29, 2018, the Commission announced a settlement with Mr. Musk and Tesla, Inc. in connection with Mr. Musk's August 7, 2018 "funding secured" statements on Twitter. The settlement imposed a \$20 million penalty on Mr. Musk and a separate \$20 million penalty on Tesla, and required Mr. Musk to step down as Chairman. The Commission's compliance enforcement of the settlement, including consent-decree obligations regarding pre-approval of certain Musk communications, extended through 2022 and beyond. A three-year CAT retention horizon would render any subsequent investigation of trading-pattern conduct in Tesla securities during 2018–2019 unreviewable by 2026.

SEC v. Billimek & Williams (Nuveen Front-Running)

On December 14, 2022, the Commission announced charges against Mr. Lawrence Billimek, a former Nuveen trader, and Mr. Alan Williams for a scheme in which Mr. Billimek tipped Mr. Williams with material non-public information about forthcoming Nuveen equity trades, generating illegal profits reported at approximately \$47 million. Several commentators characterized the matter as the Commission's first publicly credited use of Consolidated Audit Trail data to uncover a significant manipulation scheme. The case directly demonstrates CAT's value at the upper end of enforcement capability — and directly demonstrates the harm of a retention horizon that removes future comparable detections from the Commission's toolkit.

Source Materials

- U.S. Department of Justice, Press Release, "Futures Trader Charged with Illegally Manipulating Stock Market, Contributing to the May 2010 Market 'Flash Crash'" (April 21, 2015).
- Indictment, United States v. Sarao, No. 15-cr-00075 (N.D. Ill. Sept. 2, 2015).
- U.S. Securities and Exchange Commission, Press Release 2020-265, "Goldman Sachs Agrees to Pay More Than \$1 Billion to Settle SEC Charges in 1MDB Case" (October 22, 2020).
- U.S. Commodity Futures Trading Commission, Press Release 8260-20, "CFTC Orders JPMorgan to Pay Record \$920 Million for Spoofing and Manipulation" (September 29, 2020).
- U.S. Securities and Exchange Commission, Press Release 2020-233 (September 29, 2020) (J.P. Morgan Securities parallel action).

- U.S. Department of Justice, Press Release, "JPMorgan Chase & Co. Agrees to Pay \$920 Million in Connection with Schemes to Defraud Precious Metals and U.S. Treasuries Markets" (September 29, 2020).
- U.S. Securities and Exchange Commission, Press Release 2018-226, "Elon Musk Settles SEC Fraud Charges; Tesla Charged With and Resolves Securities Law Charge" (September 29, 2018).
- Public dockets and verdicts in *United States v. Hwang et al.*, S.D.N.Y.; Judge Hellerstein, sentencing transcript, November 20, 2024.
- U.S. Securities and Exchange Commission, public statements and charging documents, *SEC v. Billimek & Williams* (December 14, 2022).

PUBLIC COMMENT — FILE NO. 4-698**Exhibit C***AWS Cost-Tier Analysis: Petabyte-Level Storage Economics for CAT Data in 2026*

This exhibit documents the AWS S3 storage-class pricing used in the three-scenario cost model set forth in Section V of the accompanying public comment submission. All prices are public list prices in the US East (N. Virginia) region as of April 2026, obtained from the AWS S3 Pricing and S3 Glacier Pricing pages. Per-petabyte costs are rounded to the nearest \$500 for clarity; actual cost per byte conforms to the published list price at the fourth or fifth decimal place.

1. S3 Storage Class Pricing (April 2026, US East / N. Virginia, list prices)

Storage Class	List Price (per GB-month)	Cost per TB-year	Cost per PB-year	Retrieval Latency
S3 Standard (first 50 TB/month tier)	\$0.023	~ \$276	~ \$276,000	ms
S3 Standard-IA (Infrequent Access)	\$0.0125	~ \$150	~ \$150,000	ms
S3 Intelligent-Tiering (Frequent Access)	\$0.023	~ \$276	~ \$276,000	ms
S3 Intelligent-Tiering (Archive Access)	\$0.004	~ \$48	~ \$48,000	~5 hours
S3 Glacier Instant Retrieval	\$0.004	~ \$48	~ \$48,000	ms
S3 Glacier Flexible Retrieval	\$0.0036	~ \$43	~ \$43,200	min–5 hrs
S3 Glacier Deep Archive	\$0.00099	~ \$12	~ \$11,880	~12 hrs

Sources: aws.amazon.com/s3/pricing and aws.amazon.com/s3/glacier/pricing, retrieved April 2026. Per-TB and per-PB figures are computed as: list price (per GB-month) × 1,000 GB × 12 months (for TB-year) or × 1,000,000 GB × 12 months (for PB-year, at AWS decimal-GB billing convention), rounded. The 23× ratio between Deep Archive and Standard list pricing is reflected in the span from \$11,880 (Deep Archive) to \$276,000 (Standard) per petabyte per year.

2. Compression Arithmetic

Raw CAT records, reported in fixed-length FIX-like formats, compress aggressively when converted to columnar encoding. Apache Parquet with Zstandard compression typically delivers 4× to 8× physical compression on financial-markets order-flow data, depending on field-set cardinality and dictionary-encoding efficacy. Applied to the per-petabyte figures in Section 1 above:

- **At 4× compression**, 1 PB of raw CAT data occupies 250 TB physical; Deep Archive cost ≈ \$2,970/year.
- **At 6× compression**, 1 PB raw → 167 TB physical; Deep Archive cost ≈ \$1,980/year.
- **At 8× compression**, 1 PB raw → 125 TB physical; Deep Archive cost ≈ \$1,485/year.

The compression ratio is an engineering input independent of retention policy. Applied to the Proposed Architecture, compression savings apply on top of tier migration savings; the two levers are multiplicative, not additive.

3. Federal-Procurement Discount Range

Public federal-agency cloud-procurement data indicate that committed-use Enterprise Discount Agreements routinely yield 40% to 60% off list pricing for large-volume workloads. Additional discounts of 25% to 35% are available on committed reserved capacity with multi-year duration commitments. A compound discount of 50% on list pricing is a conservative baseline; 70% is achievable at the upper end of what the Department of Defense, the Department of the Treasury, and the National Institutes of Health have negotiated for comparable workloads.

4. Effective Per-Petabyte Economics (illustrative, combining compression and discount)

Configuration	Compression	Federal Discount	Effective PB-Year Cost (Deep Archive)
List pricing, no compression	1× (raw)	0%	~ \$11,880
Parquet + Zstd, list pricing	6×	0%	~ \$1,980
Parquet + Zstd, baseline federal discount	6×	50%	~ \$990
Parquet + Zstd, upper-end federal discount	8×	70%	~ \$446

The Proposed Architecture is materially cheaper than the approved three-year Data Storage Amendment under any reasonable combination of compression ratio and federal discount. The seven-year retention horizon is not a cost burden — it is a cost opportunity that the approved Amendment declines to capture.

5. Minimum-Storage-Duration and Retrieval Fees

Three operational notes regarding Glacier Deep Archive, drawn from the AWS service documentation, are relevant to CAT operations:

- **Minimum storage duration.** 180 days. Objects deleted before 180 days are billed for the full 180-day period. This is immaterial for CAT records intended for 7-year retention.
- **Retrieval fees.** Bulk retrieval (12-hour SLA) is billed at approximately \$0.025 per 1,000 requests plus \$0.0025 per GB restored. At 1 PB retrieved, the bulk retrieval fee is approximately \$2,500 — a rounding error against the storage cost, and fully consistent with the Commission's typical enforcement query pattern (selective subset pulls, not full-corpus restore).
- **Regulatory egress.** Standard data-transfer-out fees apply to retrieved data sent to the Internet. AWS has publicly granted zero-egress provisions for qualifying federal law-enforcement workloads; CAT regulatory pulls by the SEC and SRO Participants should be negotiated into the enterprise agreement as zero-egress.

Source Materials

- Amazon Web Services, "S3 Pricing," <https://aws.amazon.com/s3/pricing> , retrieved April 2026.
- Amazon Web Services, "S3 Glacier Pricing," <https://aws.amazon.com/s3/glacier/pricing> , retrieved April 2026.
- Amazon Web Services, "S3 Storage Classes — Glacier," <https://aws.amazon.com/s3/storage-classes/glacier> , retrieved April 2026.
- Amazon Web Services, "S3 Glacier Deep Archive Service Documentation" (minimum storage duration, retrieval fees, and Object Lock integration).
- General Services Administration, federal cloud-procurement vehicles (GSA MAS, EIS, and Alliant 2) and publicly reported Department of Defense GovCloud enterprise agreements for comparative discount benchmarking.

PUBLIC COMMENT — FILE NO. 4-698**Exhibit E***Redline of Proposed Amendments to CAT NMS Plan, Appendix D, § 6.4*

This exhibit presents the proposed textual amendments to the CAT NMS Plan, Appendix D, Sections 6.4(a), 6.4(b), and 11.1, set forth in Section IX of the accompanying public comment submission, in a formal redline format suitable for adoption by the Commission. In each redline box below, text proposed for deletion appears in strikethrough; text proposed for insertion appears in bold italics. Marginal commentary explains the rationale for each amendment.

The proposed amendments align the CAT NMS Plan's retention, privacy, and procurement provisions with (i) the Commission's statutory recordkeeping obligations under Rule 17a-4, SOX § 802, CFTC Rule 1.31, and the Bank Secrecy Act; (ii) the 10-year statute of limitations for scienter-based disgorgement codified in NDAA § 6501; and (iii) the cloud-storage economics documented in Exhibits C and D.

1. Proposed Amendment to Section 6.5(b)(i) of the CAT NMS Plan (Core Data Retention)

The Plan Processor shall retain CAT Data for a period of not less than ~~three (3) years~~ ***seven (7) years*** in respect of core Order and Execution records, and shall make such data directly available and searchable electronically without manual intervention for a period of not less than ***three (3) years***. Core Order and Execution records older than three (3) years may be migrated to an archival storage class with a retrieval service level agreement of not more than seventy-two (72) hours, provided that evidentiary integrity, immutability, and cryptographic hash verification are preserved.

Rationale

Extends the minimum retention period for core order lifecycle and execution records from three to seven years, aligning the CAT NMS Plan with: (i) the Sarbanes-Oxley § 802 audit-workpaper retention requirement (7 years); (ii) the MiFID II transaction-records regime (5-year baseline under Art. 25 MiFIR and Art. 76 of Commission Delegated Regulation 2017/565, extendable to 7 years at the request of a national competent authority; several EU Member States, including Germany, have adopted the 7-year period as the domestic default); (iii) the ten-year NDAA § 6501 statute of limitations for scienter-based disgorgement; and (iv) the empirical enforcement-timeline record set forth in Exhibit B. The three-year "hot" direct-query horizon is preserved, ensuring no degradation of routine regulatory surveillance. Records aged beyond three years are migrated to archival storage with a 72-hour retrieval SLA — wholly acceptable

for multi-year enforcement matters, trusteeship investigations, and private-plaintiff requests that are, by their nature, not time-critical.

2. Proposed Amendment to Appendix D, § 6.4(b) — PII Retention

Raw Personally Identifiable Information ("PII") fields maintained in the Customer and Account Information System shall be retained, in tokenized form, for a period of not less than *seven (7) years*. Raw, untokenized PII values shall be retained for a period of not more than *two (2) years*, after which such values shall be cryptographically destroyed while the tokenized identifiers shall continue to resolve cross-record linkages for the remainder of the seven-year retention horizon, subject to a separately secured token-mapping vault under the exclusive custody of the Plan Processor and subject to access only upon senior-officer request from the SEC or an SRO Participant.

Rationale

Bifurcates the retention horizon between tokenized identifiers (seven years, enabling cross-record linkage for enforcement purposes) and raw PII values (two years, reducing privacy-risk exposure). This architecture directly addresses the privacy objections raised in connection with the CAIS Amendment and related prior commentary. The two-year horizon for raw PII matches the shortest comparable horizon in the recordkeeping architecture (CFTC Rule 1.31's readily-accessible period); the seven-year horizon for tokenized identifiers matches SOX § 802 and MiFID II. The proposed amendment is, accordingly, more privacy-protective than the status quo with respect to raw PII while being more enforcement-protective with respect to cross-record linkage.

3. Proposed Amendment to Appendix D, § 11.1 — Cost Controls / Cloud Procurement

The Plan Processor shall implement a tiered cloud-storage lifecycle policy, including but not limited to S3 Standard, S3 Intelligent-Tiering, S3 Glacier Instant Retrieval, and S3 Glacier Deep Archive (or functionally equivalent storage classes on alternative cloud service providers), and shall *negotiate a federal Enterprise Discount Agreement and committed-use pricing with the cloud service provider, including provisions for zero-egress regulatory retrievals by the SEC and SRO Participants*. The Plan Processor shall solicit competitive bids from at least two cloud service providers meeting applicable federal security-authorization requirements (including FedRAMP High or equivalent) at intervals of not more than three (3) years.

Rationale

Codifies the tiered-storage architecture set forth in Section VI of the accompanying submission and Exhibit C, and obligates the Plan Processor to pursue federal Enterprise Discount Agreement terms and competitive

second-source bidding. The competitive-bidding requirement functions as a market-discipline mechanism analogous to that adopted by other federal agencies procuring cloud services at scale (e.g., Department of Defense, Department of the Treasury). Competitive pressure, together with tier migration and compression, is the mechanism by which the seven-year retention horizon can be achieved at lower total annual cost than the approved three-year architecture.

4. Proposed New Subsection — Appendix D, § 6.4(c): Evidentiary-Integrity Controls

*New subsection (c) to be inserted: **The Plan Processor shall apply S3 Object Lock in Compliance Mode, or a functionally equivalent immutable-storage control on any alternative cloud service provider, to CAT Data migrated to archival storage under subsection (a) of this Section. The Plan Processor shall generate and separately retain SHA-256 cryptographic hashes of every archived object, together with manifests recording the temporal and contextual metadata necessary to support chain-of-custody and authenticity challenges in subsequent enforcement or judicial proceedings.***

Rationale

Establishes evidentiary-integrity controls on archived CAT Data analogous to the WORM (Write-Once, Read-Many) storage requirements of 17 C.F.R. § 240.17a-4(f), ensuring that archival migration does not compromise the admissibility or chain-of-custody integrity of CAT records in subsequent enforcement or civil proceedings. Object Lock in Compliance Mode prevents modification or deletion of archived records during the retention period; cryptographic hashes and manifests provide tamper-evident proof of integrity that withstands authenticity challenges under Federal Rule of Evidence 901.

5. Summary of Proposed Plan Changes

The four proposed amendments above, collectively, convert the CAT NMS Plan from a three-year "delete and hope" retention posture to a seven-year tiered-archival posture that is simultaneously (i) more protective of investors, (ii) better aligned with the Commission's statutory remedies, (iii) consistent with international peer benchmarks, (iv) economically cheaper than the status quo Amendment under documented 2026 cloud pricing, and (v) drafted in a form that can be adopted by the Commission with minimal redrafting burden. The undersigned respectfully urges the Commission to adopt these amendments, or substantially equivalent amendments, in the next scheduled modification of the CAT NMS Plan.